

Empirical Evaluation of Real World Tournaments

Nicholas Mattei

*Data61/CSIRO and University of New South Wales
Sydney, Australia*

nicholas.mattei@data61.csiro.au

Toby Walsh

*Data61/CSIRO and University of New South Wales
Sydney, Australia*

toby.walsh@data61.csiro.au

Abstract

Computational Social Choice (ComSoc) is a rapidly developing field at the intersection of computer science, economics, social choice, and political science. The study of tournaments is fundamental to ComSoc and many results have been published about tournament solution sets and reasoning in tournaments [4]. Theoretical results in ComSoc tend to be worst case and tell us little about performance in practice. To this end we detail some experiments on tournaments [24] using real world data from soccer and tennis. We make three main contributions to the understanding of tournaments using real world data from English Premier League, the German Bundesliga, and the ATP World Tour: (1) we find that the NP-hard question of finding a seeding for which a given team can win a tournament is easily solvable in real world instances, (2) using detailed and principled methodology from statistical physics we show that our real world data obeys a log-normal distribution; and (3) leveraging our log-normal distribution result and using robust statistical methods, we show that the popular Condorcet Random (CR) tournament model does not generate realistic tournament data.

Keywords: Tournaments, Computational Social Choice, Economics, Preferences, Reasoning Under Uncertainty

1. Introduction

Computational Social Choice (ComSoc) has delivered impactful improvements in several real world settings ranging from optimizing kidney exchanges [14] to devising mechanisms which to assign students to schools and/or courses more fair and efficient manners [5]. ComSoc has also had impact in a number of other disciplines within computer science including recommender systems, data mining, machine learning, and preference handling [38, 7]. From its earliest days, much theoretical work in ComSoc has centered on worst case assumptions [3]. Indeed, in the last 10 years, there has been a groundswell of such research which shows little signs of slowing [11, 16, 15].

Within ComSoc, much work focuses on manipulative or strategic behavior, which may take many different forms including manipulation and control of election and aggregation functions [4]. Often, advanced algorithmic techniques such as fixed parameter tractability to move beyond these worst case assumptions [9, 16]. Approximation algorithms have played an important part, helping to determine the winner of some hard to compute voting rules [6, 40]. Approximation has been used in other areas of social choice including mechanism design, often to achieve good results when the “worst case” is too hard [36]. Additional algorithmic work has centered on average case complexity (which typically suppose very uniform sampling of instances) [35] and/or attempting to understand the parameters which make an aggregation or choice rule hard to manipulate [10, 48].

In one of the papers that founded the field, Bartholdi, III et al. [3] warned against exclusively focusing on worst case assumptions stating, “The existence of effective heuristics would weaken any practical import of our idea. It would be very interesting to find such heuristics.” For the last several years we have championed the use of real world data in ComSoc [30] and are happy to see more and more researchers working in this area (e.g., [17, 42]).¹ We see an increased focus on experimentation, heuristics, and verifying theory and models through properly incentivized data collection and experimentation as *key research direction* for ComSoc. Some of the most impactful work in ComSoc has come from the development of theory that is specifically informed by real world data and/or practical application that is then rigorously tested (e.g., [5, 14]).

Contribution. In this short note we detail a study of *tournaments* [24] using real world and generated data. We show that, despite the NP-completeness of the Tournament Fixing Problem (TFP) [2], enumerating all the seedings for which a particular player can win is quickly solvable for a large number of generated and real world instances. Additionally, we show that the popular Condorcet Random (CR) model used to generate synthetic tournaments (i) does not match real world data and (ii) is drawing from a fundamentally different distribution than real world tournaments. The statistical and modeling methodologies we use for this research may be of independent interest to empirical researchers in social choice.

2. Preliminaries: Tournaments

Whether or not teams can advance through a knockout tournament tree in order to claim a championship is a question on the minds of many during the FIFA World Cup, ATP Tennis Tournaments, NCAA Basketball Tournaments, and numerous soccer leagues around the world. The scheduling of the tournament, the seeding which dictates whom will play whom, and its manipulation in order to maximize a particular team’s chance of winning is a well studied problem in ComSoc and other areas [19, 44, 47, 20, 22, 23].

Following Aziz et al. [2], we are given a set of players $N = \{1, \dots, n\}$ and a deterministic pairwise comparisons P for all players in N . For every i, j in N , if $P_{i,j} = 1$ then

¹For a more comprehensive listing with links to over 80 papers and a wealth of tools and resources, please see www.preflib.org.

we say that player i beats player j in a head to head competition; this means that $i > j$ in a pairwise comparison. In a *balanced knockout tournament* we have that $n = 2^c$. Given a set of players N , a balanced knockout tournament $T(N, \sigma)$ is a balanced binary tree with n leaf nodes and a draw σ . There are multiple isomorphic (ordered) assignments of agents in N to the leaf nodes, we represent these as a single unordered draw σ . Observe that there are $n!$ assignments to leaf nodes but only $n!/2^{n-1}$ draws.

A knockout tournament $T(N, \sigma)$ is the selection procedure where each pair of sibling leaf nodes competes against each other. The winner of this competition proceeds up the tree into the next round; the winner of the knockout tournament is the player that reaches the root node. In this study we want to understand the computational properties of the TOURNAMENT FIXING PROBLEM (TFP).

TOURNAMENT FIXING PROBLEM (TFP):

Instance: A set of players N , a deterministic pairwise comparison matrix P , and a distinguished player $i \in N$.

Question: Does there exist a draw σ for the players in N where i is the winner of $T(N, \sigma)$?

It was recently proven that even if a manipulator knows the outcome of each pairwise matchup, the problem of finding a seeding for which a particular team will win is NP-hard [2], thus concluding a long line of inquiry into this problem. More recent results have shown that, for a number of natural model restrictions, the TFP is easily solvable [4]. Note that the code developed for the next section can be easily generalized to the case where we do not enforce the balanced constraint on T .

One popular model for generating P for experiment is the Condorcet Random (CR) model, introduced by Young [49]. The model has one tuning parameter $Pr(b)$ which gives the probability that a lower ranked team will defeat (upset) a higher ranked team in a head-to-head matchup. In general, $0.0 < Pr(b) < 0.5$. A Uniform Random Tournament has $Pr(b = 0.5)$. The manipulation of tournament seedings has been studied using both random models [41, 22, 23] and a mix of real data and random models [19]. A more sophisticated model of random tournaments was developed by Russell and van Beek [39] – upset probabilities were fitted to historical data, and varied according to the difference in the ranking of the teams (surprisingly for their tennis data, the upset probability remained relatively invariant to the difference in player rankings).

To explore how theory lines up with practice, we looked at two questions:

1. Does the NP-completeness of knockout tournament seeding manipulations tell us about the complexity of manipulation in practice?
2. Are random models which are popularly used in ComSoc supported by real world data?

To answer these questions we use data from the 2009-2014 English Premier League and German Bundesliga, along with lifetime head-to-head statistics of the top 16 tennis players on the ATP world tour to create 16 team pairwise tournament datasets.

3. Tournament Fixing in Practice

In order to convert the datasets into deterministic pairwise tournaments we used two different strategies. Using the lifetime head to head ATP tour results we get a set

Rank	Name	Seedings Won	% Total	Nodes First	Nodes All	Time All (s)
2	Novak Djokovic (SRB)	175,188,825	27.437007%	15	367	53
3	David Ferrer (ESP)	302,736	0.047413%	11	195	5.248
4	Andy Murray (GBR)	141,180,205	22.110784%	13	658	41
5	Juan Martin Del Potro (ARG)	125,622	0.019674%	11	207	2.344
6	Roger Federer (SUI)	320,366,970	50.173925%	13	4991	1
7	Tomas Berdych (CZE)	127,115	0.019908%	12	10604	2.347
8	Stanislas Wawrinka (SUI)	629	0.000099%	6	672	0.022
9	Richard Gasquet (FRA)	1,382	0.000216%	8	197	0.032
10	Jo-Wilfried Tsonga (FRA)	2,509	0.000393%	6	76	0.064
11	Milos Raonic (CAN)	1,203,771	0.188527%	23	285	21.061
12	Tommy Haas (GER)	0	0.000000%	0	N/A	N/A
13	John Isner (USA)	25	0.000004%	6	63	0.006
14	John Almagro (ESP)	0	0.000000%	0	N/A	N/A
15	Mikhail Youzhny (RUS)	13,040	0.002042%	8	398	0.241
16	Fabio Fognini (ITA)	0	0.000000%	0	N/A	N/A
17	Kei Nishikori (JPN)	46	0.000007%	6	119	0.008

Table 1: The number of knockout tournament seedings that a particular player in the world top 17 could win based on head-to-head record as of Feb. 1, 2014. Players written in **bold** are commonly called kings [31] which is equivalent to belonging to the uncovered set [24].

of weighted pairwise comparisons that we can use as input to understand tournaments. Using all data available up until Feb. 1, 2014 provides something that is not a tournament graph. There are several ties and several players who have never played each other. Given the matchup matrix, we extracted a $\{0,1\}$ tournament graph by saying one player beats another if their historical average is $\geq 50\%$. This does not create a tournament graph, hence we award all ties, including if two players have never met, to the higher ranked player. This results in Rafael Nadal being a Condorcet winner and he is thus removed from the following analysis. For the soccer data, we deterministically assigned as the pairwise winner the team which had scored more total goals in the home-and-away format for a particular year (ties broken by away goals).

We implemented a simple constraint program in miniZinc [33] to search for a seeding that ensures a given team wins. The miniZinc model which was then solved using a modified version of GeCode (<http://www.gecode.org>). The modifications to GeCode served to make the counting of all solutions faster by removing all diagnostic printing to STDOUT and implementing some other small caching optimizations. All experiments were run on a system running Debian 6.0.10 with a 2.0 GHz Intel Xeon E5405 CPU and 4 GB of RAM; GeCode was restricted to using 4 cores.

For all the tournaments in our experiments we have 16 teams which means that the entire search space is $16!/2^{15} = 638,512,875$ possible seedings. Table 1 shows the results for each of the players in the ATP World Top 16 including the total number of seedings won (and percentage of the overall total number of seedings), the number of

File	Min First	Median First	Max First	Min All	Median All	Max All
Bundesliga 2009	9	13.5	18	41	143	576
Bundesliga 2010	12	15	20	55	165	1,841
Bundesliga 2011	9	12.5	19	50	191	1,153,930
Bundesliga 2012	11	15.5	278	132	400	1,312
Bundesliga 2013	11	13.5	16	124	267	8,341
Bundesliga 2014	8	17	106	142	288	697
Premier 2009	12	15.5	22	126	295	80,497
Premier 2010	10	17	95	33	284	228,237
Premier 2011	9	12.5	63	110	358	2,823
Premier 2012	9	14.5	433	96	248	1,969
Premier 2013	6	9	15	62	364	3,016
Premier 2014	12	39	351	99	377	14,442
Tennis Top 16	6	8	23	63	285	10,604

Table 2: Summary of the number of choice points explored to find the first or all seedings for which a team can win a knockout tournament, respectively, for all teams across all datasets. Results from finding seedings for which any given team can win a knockout tournament. Nodes are the number of choice points explored to find the first or all seedings, respectively, for all teams. Teams which could not win any seeding are not included in these results.

choice points (nodes of the search tree) explored to find the first (resp. all) seedings, and and time (minutes) explored to to find all seedings for which a player can win a tournament. Table 2 gives summary statistics for the number of choice points (nodes of the search tree) and time (minutes) explored to to find the first (resp. all) seedings for which a team can win a knockout tournament for all teams across all 13 datasets.

Despite this being a NP-hard problem, it was extremely easy in every dataset to find a winning seeding for any given team (or prove that none existed). Exhaustively, counting all the winning seedings took more time but even this was achievable. Only the Bundesliga 2011 experiment came anywhere close to exploring even a small fraction of the $\frac{16!}{2^{15}} = 638,512,875$ total possible seedings.

The practical upshot of these computational results is that real world tournaments exhibit a lot of structure which is possible to leverage for practical computation. While the simple techniques we employed may not scale to the $\frac{128!}{2^{127}} = 2.3 * 10^{177}$ possible seedings in the 128 player Wimbledon Tournament, they can be used for more modestly sized tournaments. The low number of choice points explored for these instances may indicate that there is practically exploitable structure in larger tournaments; an interesting avenue for future research.

4. Verifying Real World Tournament Models

We turn to a second fundamental question. The CR model has been used to derive theoretical bounds on the complexity of manipulating the seeding in knockout tourna-

ments [41]. But does it adequately model real world tournaments? In the soccer datasets we take the teams goals scored over the total as the pairwise probability, while we use the life time head-to-head average to determine probabilities in the tennis data. To test the modeling power of CR, we generated pairwise probabilities with $0.0 < Pr(b) \leq 0.5$. We then computed, for each of the real world datasets and all values of $Pr(b)$ in steps of 0.01, the probabilities that teams would win a knockout tournament using a simple sampling procedure which converged to the actual probability quickly; uniformly sampling over all possible seedings and updating the probability estimates for each team for a particular seeding, the probability that every team wins the tournament can be computed efficiently [29, 46, 45].

Our hypothesis is that the probability of winning a knockout tournament for a team in the real world data is a random variable drawn from the same distribution as the CR model. We approach this question in two parts. (1) Using a Kolmogorov-Smirnov (KS) tests with $p = 0.05$ as our significance threshold [8], we can determine if the data is drawn from the same type of distribution, i.e. a normal distribution or a heavy tailed distribution such as a log-normal or power law distribution. (2) Then for a candidate pair of samples, we determine if the fitting parameters of the distribution are similar.

The KS test compares the distance between the cumulative distribution (CDF) of two empirical samples to determine whether or not two different samples are drawn from the same distribution. Figure 1(A) shows the CDF of the 2014 Bundesliga league data along with several settings of $Pr(b)$. Table 3 gives the minimum and maximum values of $Pr(b)$, per dataset, for which we can say that the probability distribution of a team winning a knockout tournament according to the CR model is likely drawn from the same distribution as the respective real world dataset (KS test, $p = 0.05$). We can reject CR models with values of $Pr(b)$ outside these ranges; as these models are not likely to emerge from the same distribution as the real world datasets. We also provide average upset probability for each datafile to compare with the results of Russell and van Beek [39].

Examining our results, we find no support for the Uniform Random Tournament model. Likewise, setting $Pr(b) < \approx 0.13$ or $Pr(b) > \approx 0.42$ generates data which is drawn from a different distribution than most real world datasets we survey. The tennis data seems to be an outlier here, supporting a very low value of $Pr(b)$, likely due to Rafael Nadal, who has a winning lifetime record against all other players in the ATP top 16 as of Feb. 1, 2014.

As we cannot reject all models given by CR outright we must look more closely at the underlying distribution and attempt to fit the empirical data to a likely distribution. For this we will dive more deeply into the 2014 Bundesliga League data, as the range for $Pr(b)$ is similar to the average of $0.153 < Pr(b) < 0.42$ across all datasets and the average upset probability for 2014 yields a model which is a good match for the underlying data. The 2014 Bundesliga data has an average upset probability of $Pr(b = 0.374)$ and a best fit probability according to the KS test of $Pr(b = 0.30)$.


We must first identify what kind of probability distribution the samples are drawn from in order to tell if they are the same or different. At first glance, the winning probabilities appear to be drawn from a power law or some other heavy tailed distribution such as a log-normal [32, 8, 25]. The study of heavy tailed distributions in empirical data is a rich topic that touches a number of disciplines including physics,

Data	Min $Pr(b)$	Max $Pr(b)$	Average Upset Probability
Bundesliga 2009	0.13	0.47	0.38433
Bundesliga 2010	0.15	0.45	0.39714
Bundesliga 2011	0.21	0.45	0.41114
Bundesliga 2012	0.20	0.42	0.38266
Bundesliga 2013	0.12	0.43	0.37767
Bundesliga 2014	0.15	0.39	0.37401
Premier 2009	0.12	0.43	0.34417
Premier 2010	0.14	0.37	0.33683
Premier 2011	0.17	0.44	0.41291
Premier 2012	0.15	0.39	0.39523
Premier 2013	0.16	0.41	0.40010
Premier 2014	0.13	0.39	0.33184
Tennis Top 16	0.04	0.35	0.29961


Table 3: Minimum and Maximum of the range of $Pr(b)$ for which we can say that the probability distribution of a team winning a knockout tournament according to the CR model is drawn from the same distribution as the respective real world dataset (KS test, $p = 0.05$). We also provide average upset ($Pr(b)$) probability for each datafile.

computer science, literature, transportation science, geology, biology as these distributions describe a number of natural phenomena such as the spread of diseases, the association of nodes in scale free networks, the connections of neurons in the brain, the distribution of wealth amongst citizens, city sizes, and other interesting phenomena [1, 8, 25]. Recently, more sophisticated methods of determining if an empirical distribution follows a particular heavy tailed distribution have been developed, consequently showing strong evidence that distributions once thought power laws (e.g., node connections on the internet and wealth distribution) are likely not explained by a power law distribution [8] but rather by log-normal distributions [25]. The current standard for fitting heavy tailed distributions in physics and other fields (and the one we will employ) involves the use of robust statistical packages to estimate the fitting parameters then testing the fitted model for basic plausibility through the use of a likelihood ratio test [1]. This process will help us decide which distribution is the strongest fit for our data as well as provide us with the actual fitting parameters to compare the real world and generated data.

Figure 1 (B) shows the results of fitting the 2014 Bundesliga League data to a power law for a random variable X of the form $Pr(X \geq x) \propto cx^{-\alpha}$ as well as the fit for a log-normal distribution $Pr(X \geq x) \propto \Delta(\mu, \sigma)$ with median μ and multiplicative standard deviation σ . Using a likelihood ratio test we find that the log-normal is a significantly better fit for the data than the power law distribution ($R = 0.7319$, $p = 0.4642$). This makes intuitive sense in this context as each matchup can be seen as a (somewhat) independent random variable, and the product of multiple positive random variables gives a log-normal distribution [25]. The fit parameters for the 2014 Bundesliga League data are $\sigma = 1.2611$ and $\mu = -4.0717$ while for the best fitting CR model with $Pr(b =$



(A)



(B)

Figure 1: (A) Cumulative distribution function (CDF) of the probability of a team winning a tournament of the 2014 Bundesliga League data along with several random benchmarks. *Average* is the value of $Pr(b)$ computed as the average in the dataset while *Best* is the value of $Pr(b)$ tested (in 0.01 increments) which minimizing the KS distance. (B) Comparison of fitted probability distributions for the 2014 Bundesliga League data. The Log-Normal distribution is the best fit according to the likelihood ratio test.

0.30) is $\sigma = 1.0018$ and $\mu = -3.3823$. While those two distributions are similar, it implies that perhaps a more nuanced, multi-parameter model is needed to capture the matchup probabilities for tournaments.

5. Discussion and Future Directions

In order to transform into a more impactful area we must demonstrate the effectiveness of our methods in real settings, and let these real settings drive our research. Kagel and Roth [21] describe the journey for experimental economics: evolving from theory to simulated or repurposed data to full fledged laboratory and field experiments. This progression enabled a “conversation” between the experimentalists and the theoreticians which enabled the field to expand, evolve, and have the impact that it does today in a variety of contexts.

Our case study of tournaments shows that the need to verify our models with data is an important and interesting future direction for ComSoc. Working to verify these models can point the way to new domain restrictions or necessary model generalizations. There has been more data driven research, like the kind presented here, thanks to PrefLib [30] and other initiatives, e.g., [17, 42]; we hope this trend continues. This research complements the existing research on axiomatic characterizations, worst case complexity, and algorithmic considerations (see, e.g., [13, 40, 18]). While there are

some issues with using repurposed data and discarding context (see, e.g., the discussion by John Langford of Microsoft Research about the UCI Machine Learning Research Repository at <http://hunch.net/?p=159>) it is a *start* towards a more nuanced discussion about mechanisms and preferences.

Results about preference or domain restrictions can lead to elegant algorithmic results, but these restrictions should be complemented by some evidence that the restrictions are applicable. For example, in the over 300 complete, strict order datasets in PrefLib [30], none are single-peaked, a popular profile restriction in voting [16]. Untested normative or generative data models can sometimes lead us astray; if we use self reported data from surveys, or hold out data which does not fit our preconceived model, we may introduce bias and thus draw conclusions that are spurious [37, 34]. Our study of tournaments makes unrealistic assumptions about model completion in order to produce deterministic tournaments. However, even these simple rounding rules yield instances that are much simpler than the worst case results would imply.

Perhaps one way forward for ComSoc is incorporating even more ideas from experimental economics [21] including the use of human subjects experiments on Mechanical Turk [28, 12] like those performed by, e.g., Mao and Suri [27] and Mao et al. [26]. Work with human subjects can lead to a more refined view of strategic behavior and inform more interesting and realistic models on which we can base good, tested theory [43, 42].

Acknowledgments. We would like to thank Haris Aziz for help collecting data and insightful comments. Data61/CSIRO (formerly known as NICTA) is funded by the Australian Government through the Department of Communications and the ARC through the ICT Centre of Excellence Program.

References

- [1] Alstott, J., Bullmore, E., Plenz, D., 2014. powerlaw: A Python package for analysis of heavy-tailed distributions. *PloS one* 9 (1), e85777.
- [2] Aziz, H., Gaspers, S., Mackenzie, S., Mattei, N., Stursberg, P., Walsh, T., 2014. Fixing a balanced knockout tournament. In: *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*.
- [3] Bartholdi, III, J., Tovey, C., Trick, M., 1989. The computational difficulty of manipulating an election. *Social Choice and Welfare* 6 (3), 227–241.
- [4] Brandt, F., Conitzer, V., Endriss, U., Lang, J., Procaccia, A. D. (Eds.), 2016. *Handbook of Computational Social Choice*. Cambridge University Press.
- [5] Budish, E., Cantillon, E., 2012. The multi-unit assignment problem: Theory and evidence from course allocation at Harvard. *The American Economic Review* 102 (5), 2237–2271.
- [6] Caragiannis, I., Kaklamanis, C., Karanikolas, N., Procaccia, A. D., 2014. Socially desirable approximations for Dodgson’s voting rule. *ACM Transactions on Algorithms (TALG)* 10 (2), 6.

- [7] Chevaleyre, Y., Endriss, U., Lang, J., Maudet, N., 2008. Preference handling in combinatorial domains: From AI to social choice. *AI Magazine* 29 (4), 37–46.
- [8] Clauset, A., Shalizi, C. R., Newman, M. J., 2009. Power-law distributions in empirical data. *SIAM review* 51 (4), 661–703.
- [9] Conitzer, V., 2010. Making decisions based on the preferences of multiple agents. *Communications of the ACM* 53 (3), 84–94.
- [10] Conitzer, V., Sandholm, T., 2006. Nonexistence of voting rules that are usually hard to manipulate. In: *Proceedings of the 21st AAAI Conference on Artificial Intelligence (AAAI)*. pp. 627–634.
- [11] Conitzer, V., Sandholm, T., Lang, J., 2007. When are elections with few candidates hard to manipulate? *Journal of the ACM* 54 (3), 14.
- [12] Crump, M. J. C., McDonnell, J. V., Gureckis, T. M., 2013. Evaluating Amazon’s Mechanical Turk as a tool for experimental behavioral research. *PloS one* 8 (3), e57410.
- [13] Davies, J., Katsirelos, G., Narodytska, N., Walsh, T., 2011. Complexity of and algorithms for Borda manipulation. In: *Proceedings of the 25th AAAI Conference on Artificial Intelligence (AAAI)*. pp. 657–662.
- [14] Dickerson, J. P., Procaccia, A. D., Sandholm, T., 2012. Optimizing kidney exchange with transplant chains: Theory and reality. In: *Proceedings of the 11th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*. pp. 711–718.
- [15] Faliszewski, P., Hemaspaandra, E., Hemaspaandra, L. A., 2010. Using complexity to protect elections. *Journal of the ACM* 53 (11), 74–82.
- [16] Faliszewski, P., Procaccia, A. D., 2010. AI’s war on manipulation: Are we winning? *AI Magazine* 31 (4), 53–64.
- [17] Goldman, J., Procaccia, A. D., 2014. Spliddit: Unleashing fair division algorithms. *jacm* 13 (2), 41–46.
- [18] Goldsmith, J., Lang, J., Mattei, N., Perny, P., 2014. Voting with rank dependent scoring rules. In: *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*.
- [19] Hazon, N., Dunne, P. E., Kraus, S., Wooldridge, M., 2008. How to rig elections and competitions. In: *Proceedings of the 2nd International Workshop on Computational Social Choice (COMSOC)*.
- [20] Horen, J., Riezman, R., 1985. Comparing draws for single elimination tournaments. *Operations Research* 33 (2), 249–262.
- [21] Kagel, J. H., Roth, A. E., 1995. *The Handbook of Experimental Economics*. Princeton University.

- [22] Kim, M. P., Sukhompong, W., Vassilevska Williams, V., 2016. Who can win a single-elimination tournament? In: Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI). pp. 516–522.
- [23] Kim, M. P., Vassilevska Williams, V., 2015. Fixing tournaments for kings, chokers, and more. In: Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI). pp. 561–567.
- [24] Laslier, J. F., 1997. Tournament Solutions and Majority Voting. Springer-Verlag, Berlin.
- [25] Limpert, E., Stahel, W. A., Abbt, M., 2001. Log-normal distributions across the sciences: Keys and clues. *BioScience* 51 (5).
- [26] Mao, A., Procaccia, A. D., Chen, Y., 2013. Better human computation through principled voting. In: Proceedings of the 27th AAAI Conference on Artificial Intelligence (AAAI).
- [27] Mao, A., Suri, S., 2013. How, when, and why to do online behavioral experiments. In: Proceedings of the 9th International Workshop on Internet and Network Economics (WINE). Tutorial Forum.
- [28] Mason, W., Suri, S., 2012. Conducting behavioral research on Amazon’s Mechanical Turk. *Behavior Research Methods* 44 (1), 1–23.
- [29] Mattei, N., Goldsmith, J., Klapper, A., Mundhenk, M., 2015. On the complexity of bribery and manipulation in tournaments with uncertain information. *Journal of Applied Logic* 13 (4), 557–581.
- [30] Mattei, N., Walsh, T., 2013. Preflib: A library for preferences, [HTTP://WWW.PREFLIB.ORG](http://www.preflib.org). In: Proceedings of the 3rd International Conference on Algorithmic Decision Theory (ADT).
- [31] Maurer, S. B., 1980. The king chicken theorems. *Mathematics Magazine*, 67–80.
- [32] Mitzenmacher, M., 2004. A brief history of generative models for power law and lognormal distributions. *Internet Mathematics* 1 (2), 226–251.
- [33] Nethercote, N., Stuckey, P. J., Becket, R., Brand, S., Duck, G. J., Tack, G., 2007. Minizinc: Towards a standard cp modelling language. In: Proceedings of the 13th International Conference on Principles and Practice of Constraint Programming (CP). Springer, pp. 529–543.
- [34] Popova, A., Regenwetter, M., Mattei, N., 2013. A behavioral perspective on social choice. *Annals of Mathematics and Artificial Intelligence* 68 (1–3), 135–160.
- [35] Procaccia, A. D., Rosenschein, J. S., 2007. Junta distributions and the average-case complexity of manipulating elections. *Journal of Artificial Intelligence Research* 28, 157–181.

- [36] Procaccia, A. D., Tennenholtz, M., 2009. Approximate mechanism design without money. In: Proceedings of the 10th ACM Conference on Electronic Commerce (ACM-EC). pp. 177–186.
- [37] Regenwetter, M., Grofman, B., Marley, A. A. J., Tsetlin, I. M., 2006. Behavioral Social Choice: Probabilistic Models, Statistical Inference, and Applications. Cambridge Univ. Press.
- [38] Rossi, F., Venable, K., Walsh, T., 2011. A Short Introduction to Preferences: Between Artificial Intelligence and Social Choice. *Synthesis Lectures on Artificial Intelligence and Machine Learning* 5 (4), 1–102.
- [39] Russell, T., van Beek, P., 2011. An empirical study of seeding manipulations and their prevention. In: Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI). pp. 350–356.
- [40] Skowron, P., Faliszewski, P., Slinko, A., 2013. Achieving fully proportional representation is easy in practice. In: Proceedings of the 12th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS). pp. 399–406.
- [41] Stanton, I., Vassilevska Williams, V., 2011. Manipulating stochastically generated single-elimination tournaments for nearly all players. In: Proceedings of the 7th International Workshop on Internet and Network Economics (WINE). pp. 326–337.
- [42] Tal, M., Meir, R., Gal, Y., 2015. A study of human behavior in online voting. In: Proceedings of the 14th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS). pp. 665–673.
- [43] Thompson, D. R. M., Lev, O., Leyton-Brown, K., Rosenschein, J., 2013. Empirical analysis of plurality election equilibria. In: Proceedings of the 12th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS). pp. 391–398.
- [44] Vassilevska Williams, V., 2010. Fixing a tournament. In: Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI). AAAI Press, pp. 895–900.
- [45] Vu, T., Altman, A., Shoham, Y., 2009. On the complexity of schedule control problems for knockout tournaments. In: Proceedings of the 8th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS). pp. 225–232.
- [46] Vu, T., Hazon, N., Altman, A., Kraus, S., Shoham, Y., Wooldridge, M., 2013. On the complexity of schedule control problems for knock-out tournaments.
- [47] Vu, T., Shoham, Y., 2011. Fair seeding in knockout tournaments. *ACM Transactions on Intelligent Systems Technology* 3 (1), 1–17.

- [48] Xia, L., Conitzer, V., 2008. Generalized scoring rules and the frequency of coalitional manipulability. In: Proceedings of the 9th ACM Conference on Electronic Commerce (ACM-EC). pp. 109–118.
- [49] Young, H. P., 1988. Condorcet’s theory of voting. *The American Political Science Review* 82 (4), 1231 – 1244.